

Automated morphological classification of APM galaxies by supervised artificial neural networks

A. Naim,¹ O. Lahav,¹ L. Sodré, Jr² and M. C. Storrie-Lombardi¹

¹*Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA*

²*Instituto Astronômico e Geofísico da Universidade de São Paulo, CP9638, 01065-970, São Paulo, Brazil*

Accepted 1995 January 31. Received 1995 January 27; in original form 1994 November 3

ABSTRACT

We train artificial neural networks to classify galaxies based solely on the morphology of the galaxy images as they appear on blue survey plates. The images are reduced, and morphological features such as bulge size and the number of arms are extracted, all in a fully automated manner. The galaxy sample was first classified by six independent experts. We use several definitions for the mean type of each galaxy, based on those classifications. We then train and test the network on these features. We find that the rms error of the network classifications, as compared with the mean types of the expert classifications, is 1.8 Revised Hubble types. This is comparable to the overall rms dispersion between the experts. This result is robust and almost completely independent of the network architecture used.

Key words: catalogues – galaxies: general.

1 INTRODUCTION

Since the introduction of the Hubble classification scheme (Hubble 1926, 1936), astronomers have been looking at ways to classify galaxies. Other systems were suggested, e.g. Mt. Wilson (Sandage 1961), Yerkes (Morgan 1958), Revised Hubble (de Vaucouleurs 1959) and DDO (van den Bergh 1960a, b, 1976); each has its special characteristics, but they all share Hubble's original notion that the sequence of morphologies attests to an underlying sequence of physical processes.

This notion has been widely accepted for the past few decades, making morphological classification of large numbers of galaxies important for better modelling and understanding of galaxy structure and evolution. Examples include statistical relations which are specific to certain types of galaxies, e.g. the D_n – σ relation for ellipticals (Lynden-Bell et al. 1988), the Tully–Fisher relation for spirals (Tully & Fisher 1977) and the morphology–density relation (Hubble 1936; Dressler 1980).

Morphological classification of galaxies is usually done by visual inspection of photographic plates. This is by no means an easy task, requiring skill and experience. It is also time-consuming: catalogues containing human classifications take years to complete and contain of order 10^4 entries [e.g. the Third Reference Catalogue of Bright Galaxies (de Vaucouleurs et al. 1991) and the ESO catalogue (Lauberts & Valentijn 1989)]. However, in the APM (Automated Plate

Measuring machine) survey (e.g. Maddox et al. 1990) there are roughly 2×10^6 galaxies, and the expected yield of the Sloan Digital Sky Survey (Gunn et al., in preparation) is over 10^7 CCD images of galaxies. Clearly, such numbers of galaxies cannot be classified by humans. There is an obvious need for automated methods that will put the knowledge and experience of the human experts to use and produce very large samples of automatically classified galaxies.

The first stage towards achieving this goal was creating a uniform, well-defined sample to be classified by human experts. This was done in previous papers (Lahav et al. 1994; Naim et al. 1995, hereafter Paper I), where the same sample of galaxies was presented to six independent expert observers and a detailed analysis of their classifications was carried out. The experts are R. Buta, H. Corwin, G. de Vaucouleurs, A. Dressler, J. Huchra and S. van den Bergh (hereafter RB, HC, GV, AD, JH and vdB, respectively). We found that the rms dispersions between pairs of experts ranged from 1.3 to 2.1 Revised Hubble types, and that the overall rms dispersion was 1.8 types.

The next stage, which is carried out in this paper, entails training a computer software to classify galaxies on the basis of their apparent morphology. Our choice of an automated classifier is artificial neural networks (ANNs), which proved in a pilot study (Storrie-Lombardi et al. 1992) to be well suited for this task. The original idea behind ANNs was the creation of a simplified model of the human brain (McCulloch & Pitts 1943; Hopfield & Tank 1986), but they

were found to be well suited for a variety of applications in astronomy, such as classifying objects in the *IRAS* point source catalogue (e.g. Adorf & Meurs 1988), adaptive optics (e.g. Angel et al. 1990), scheduling observation time (e.g. Adorf 1989) and star–galaxy separation (e.g. Odewahn et al. 1991).

In this stage of our research we apply ‘Supervised Learning’, whereby we attempt to teach the ANN to mimic the human classifications. The ANN is given a set of parameters describing each galaxy and is told what the ‘correct’ type is. It then tries to make its classification as similar to the desired one as possible. (For more details on supervised learning, as well as other aspects of neural networks in the context of galaxy classification, see Lahav et al. 1995, in preparation.)

The outline of this paper is as follows. In Section 2 we discuss the galaxy sample. The various ways of defining mean galaxy types for training the ANN are explained in Section 3. In Section 4 we describe the preparation of galaxy images for feature extraction. The process of extracting morphological features for the ANN is described in detail in Section 5. In Section 6 we give the results of training various configurations of the ANN, based on different choices of input parameters, mean types and ANN architectures. The discussion follows in Section 7. In the Appendix we give a detailed listing of the human and ANN classifications, as well as of all the parameters the ANNs used, for a portion of the whole sample. The full table for the entire sample may be obtained from the authors.

2 THE GALAXY SAMPLE

The galaxies were all taken from the APM Equatorial Catalogue of Galaxies (Raychaudhury et al., in preparation), which is 98 per cent complete for galaxies of magnitude $B \leq 17$ mag and $D \geq 0.5$ arcmin, covering most of the sky in the ranges of declination $-17.5^\circ < \delta < 2.5^\circ$, and galactic latitude $b \geq 20^\circ$. The plates were IIIa-J (broad blue–green band) plates taken with the 48-inch UK Schmidt telescope at Siding Spring, Australia. All the images were scanned from glass copies of the original plates with a resolution of 1 arcsec by the APM machine. The plates themselves, however, have a resolution of roughly 2 arcsec (due to observing conditions), which is therefore the limiting resolution of the digitized images. No plate matching was performed, and no account was taken of possible brightness gradients within plates. Calibration for the plates was not available at the time the sample was compiled, but it later turned out that the rms difference between the plates (~ 0.2 mag) was not large enough to become crucial. At any rate, the images had already been classified by the experts (without plate matching) at that stage, and for the sake of consistency we used the same images for the ANN.

A diameter-limited sample ($D > 1.2$ arcmin at an isophotal level of 24.5 mag arcsec $^{-2}$) from 75 plates was compiled. The APM machine scans plates in strips 2.1-arcmin wide,

and at the time this compilation was made the strips were analysed separately, so large images were sometimes broken down (depending on their orientation with respect to the scanning strip). For this reason the original list of APM-selected galaxies with $D > 1.2$ arcmin was augmented by galaxies from the PGC catalogue (Paturel et al. 1989) with $D > 0.9$ arcmin. The resulting sample contained 835 galaxy images, which were then sent for classifying by the human experts (see Paper I for full details). 830 of these images were subsequently used for training the ANN (five were rejected due either to low picture quality or to being duplicates of other images already in the sample).

3 MORPHOLOGICAL TYPES

The full description of the expert classifications appears in Paper I. For training the ANN, only *T*-types (Table 1) in the range -5 to 10 were considered as classifications. During the training phase the ANN may be presented with all the different classifications each galaxy received by the experts (see below for further details), or it may be given just one mean classification. In any case, since the ANN operates as a function from the inputs to the output, it can produce only one result. Therefore, regardless of the way we present the inputs to the ANN, it must be provided with one desired output value per galaxy (the ‘correct’ answer). We considered the following options for defining the desired outputs:

- (i) a straight unweighted mean of the types given by all experts;
- (ii) a corrected mean type that takes little account of outlying classifications;
- (iii) a similarly corrected mean of a subset of experts, who agree with each other better than with the others, and
- (iv) the classifications of a single expert.

For option (i) the straight mean was calculated for each galaxy as a single number, and the resulting list of numbers was taken to be ‘type set 1’.

Option (ii) requires some algorithm for disregarding outlying classifications. We chose an iterative scheme of weighted means. The straight mean is used as the value of iteration zero, and for each galaxy the new mean type is calculated as

$$\langle T_{i+1} \rangle = \sum_{j=1}^{N_{\text{exp}}} \frac{T^j}{|T^j - \langle T_i \rangle| + 1} \bigg/ \sum_{j=1}^{N_{\text{exp}}} \frac{1}{|T^j - \langle T_i \rangle| + 1}, \quad (1)$$

where $\langle T_i \rangle$ is the mean type in iteration i and T^j is the type given to the current galaxy by expert j . For each galaxy the summations are only over those experts who actually gave it a *T*-type. The effect of this weighting scheme is to give lower weights to classifications that are further away from the previous estimate of the mean type. Successive iterations drive the mean closer and closer to the mean of the ‘good’ points. We carried out several such iterations, and Fig. 1

Table 1. *T*-types in the Revised Hubble System. For training the ANN only types in the range $[-5, 10]$ were considered.

<i>cE</i>	<i>E0</i>	<i>E+</i>	<i>S0-</i>	<i>S0°</i>	<i>S0+</i>	<i>S0/a</i>	<i>Sa</i>	<i>Sab</i>	<i>Sb</i>	<i>Sbc</i>	<i>Sc</i>	<i>Scd</i>	<i>Sd</i>	<i>Sdm</i>	<i>Sm</i>	<i>Im</i>	<i>cI</i>	<i>I0</i>	<i>Pec</i>
-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8	9	10	11	90	99

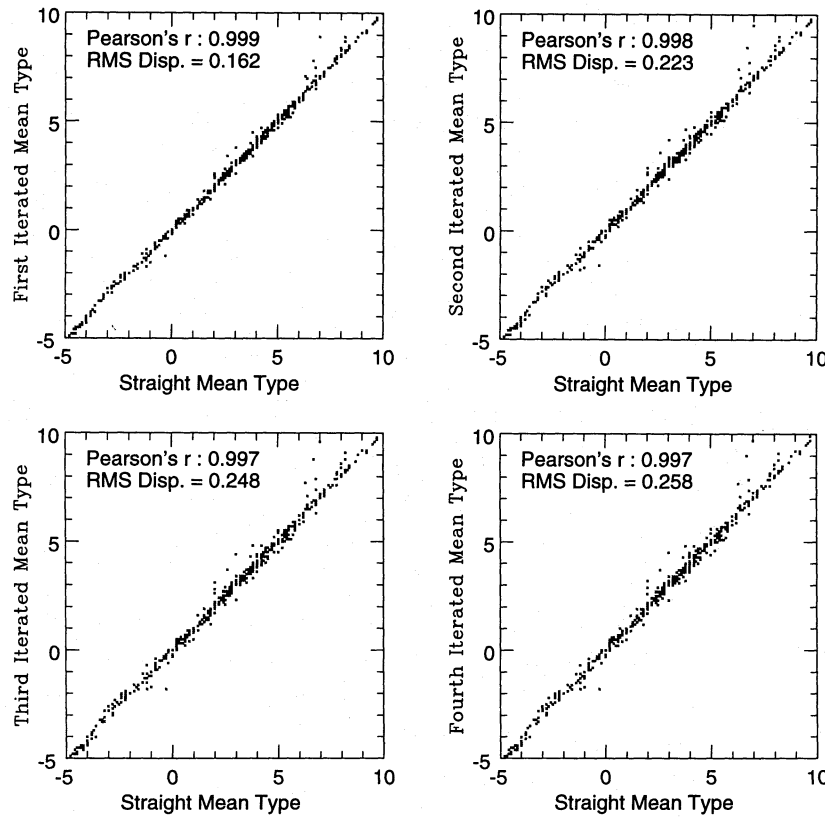


Figure 1. Iterative weighted mean type for all six experts. The plots show for each galaxy the current iteration mean type versus the original straight mean type. The correlation between the straight mean type and the iterated mean type becomes worse as the outliers become less important in successive iterations. Root mean square values and Pearson's r are calculated between the current iteration mean type and the straight mean.

shows the scatter plots of the iterated means versus the straight mean. The change in the rms value becomes smaller from one iteration to the next, and after four iterations it becomes small enough for us to stop there. We took as 'type set 2' the means arrived at after four successive iterations. As will be explained below (Section 6.2.2), in some cases we presented the ANN with all the available classifications for each galaxy (i.e. the same galaxy would appear several times, each time with a classification given by a different expert), and asked it to arrive as close as possible to the corrected mean types of type set 2.

For option (iii) we chose the classifications of RB, HC and GV, for whom we had found (in Paper I) the tightest correlations. Again we employed the approach of the iterative weighted mean, and chose as 'type set 3' the means arrived at after four successive iterations. Scatter plots for this case are shown in Fig. 2.

Finally, for option (iv) we took RB's classifications to make 'type set 4', HC's for 'type set 5', GV's for 'type set 6', AD's for 'type set 7', JH's for 'type set 8', and vdB's for 'type set 9'. Table 2 summarizes the choice of type sets.

4 IMAGE REDUCTION

The original images were kept in FITS format files of either 256×256 or 512×512 pixel. Pixel values represent density measurements from the APM microdensitometer. Several stages were required in order to standardize the images:

- (i) subtracting the sky background;
- (ii) removing superposed foreground stars, and
- (iii) standardization of images in size.

These stages were carried out automatically by software written for this purpose.

4.1 Sky subtraction

A heavily smoothed version of each image was prepared and a histogram of pixel intensities was made for it. Normally the sky background would stand out as the most prominent (and lowest in pixel intensity) peak in the histogram. A Gaussian would be fitted to that peak and a threshold value set to 2σ above the peak, in terms of the unsmoothed image. (Smoothing sharpens the peak considerably so that the width of the fitted Gaussian cannot be taken as the real σ .) However, in 11 cases the galaxy was too faint and was mistaken by the automatic program to be part of the sky. In these cases an artificial sky level was fed in by hand. Whichever way the threshold was determined, once set it was subtracted from the whole image and each pixel resulting with a negative value was set to zero. An estimate of the galaxy size was made and any remaining features outside it (e.g. stellar images) were set to zero as well. Thus at the end of this stage the only remaining stellar images were superposed on the galaxy image.

The automatic procedure described above results in some inaccuracies in the evaluation of the exact sky level to be

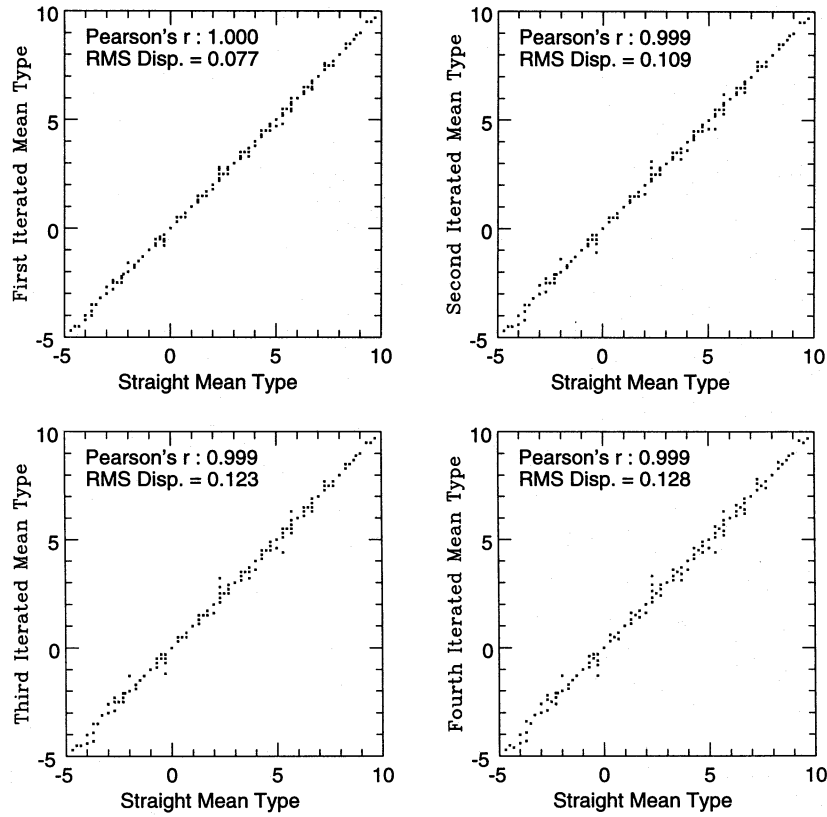


Figure 2. Iterative weighted mean type for RB, HC and GV. The plots show for each galaxy the current iteration mean type versus the original straight mean type. Root mean square values and Pearson's r are calculated between the current iteration mean type and the straight mean.

Table 2. Choices of type sets.

Type Set number	Based on Expert(s)	No. of Galaxies	Remarks
1	RB,HC,GV,AD,JH,vdB	830	straight (unweighted) mean
2	RB,HC,GV,AD,JH,vdB	830	iteratively corrected mean
3	RB,HC,GV	830	iteratively corrected mean
4	RB	764	
5	HC	812	
6	GV	473	
7	AD	814	
8	JH	824	
9	vdB	549	

subtracted. However, in most cases this inaccuracy is small enough to be insignificant for our purposes. We are not doing photometry here, but rather considering morphological features of an entire image, and these are not very sensitive to the sky threshold we calculate.

4.2 Star removal

For each image a list of star candidates was prepared by binning all intensities into a few levels and picking up local maxima. The exact number of bins was dictated by the overall variability of the image. Of these candidates the ones which qualified as stars had to have a nearly Gaussian light profile. Intensities were measured in circular annuli around the centre of each star candidate. Allowing for the potentially burnt-out centre meant that Gaussian fitting could not

proceed all the way to the centre of the stellar image, while, since the star was embedded in patches of the galaxy, it was difficult to go very far out. Typically this left us with only 3–4 annuli that we could count on, and for this reason we checked only the radii for which the intensity first dropped below 80 and 60 per cent of the intensity at the centre of the star candidate. If the profile is indeed Gaussian, the intensities at these radii are related through

$$\frac{I(r_{80})}{I(r_{60})} = \exp\left(-\frac{r_{80}^2 - r_{60}^2}{2\sigma^2}\right) \approx \exp\left(-\frac{r_{80}^2}{2r_{60}^2} + \frac{1}{2}\right), \quad (2)$$

where $I(r)$ is the intensity at a given radius r , and r_{80} and r_{60} are the radii of 80 and 60 per cent central intensity, respectively. In equation (2) use has been made of the approximation $r_{60} \approx \sigma$. Since $I(r_{80})/I(r_{60})$ is by definition 0.8/0.6, a

Gaussian would give $r_{80}/r_{60}=0.65$. Since the resolution is low, both r_{80} and r_{60} are highly discretized, and in order to overcome this problem we adopted the condition $r_{80}/r_{60} \geq 0.65$ as a criterion for a stellar image. Further criteria were employed as well. To avoid removing star-like H II regions, only very bright star candidates (those whose maximum was in the top three bins) were considered for removal; also, to avoid mistaking very bright and nearly Gaussian parts of the galaxy for stars, a maximum stellar image size was decided, and no candidate larger than this size was removed.

For each removed star a radius of deletion was set at the annulus where the average intensity dropped for the last time, before beginning to rise again. The area to be removed was set to intensity zero and then the gap was filled one annulus at a time from the outside in, each new pixel taking on the average values of its nearest non-zero neighbours. This worked well automatically for most cases. Nevertheless, there were 15 cases in which superposed stars were too big

compared to the image of the galaxy to be correctly removed by the program, and a list of stars for deletion was prepared by hand for these galaxies. In some other galaxies bright stars were superposed on very bright patches of the galaxy and were not identified as stars at all. However, the vast majority of superposed stars were removed successfully. An example of one galaxy image before and after sky subtraction and star removal is shown in the top part of Fig. 3. The processed image has been slightly smoothed to reduce the effect of noise in the feature extraction stage.

4.3 Image standardization

The centre of each galaxy was located and translated to the centre of the pixel map. Second moments were calculated on the entire image in order to find its ellipticity and position angle. The ellipticity was defined as $1 - b/a$, where b is the semiminor axis and a the semimajor axis. The position angle




Figure 3. Top: f634g002, before (left) and after (right) sky subtraction and star removal. The experts were sent the original images, which looked like the one on the left. Bottom: f646g004: sky-subtracted and star-removed (left) versus reconstructed from the elliptical annuli (right). The reconstructed image was enlarged to show its lower resolution. It is also slightly rotated compared to the original image.

was measured counterclockwise from the positive x -axis of the picture. The size of the galaxy was also estimated and its surface brightness worked out as the average pixel intensity over the entire image.

The amount of information (i.e. the number of pixels) available varies from image to image, as a function of diameter and inclination. In order to get a standard number of picture elements from all galaxies, regardless of size and inclination, the images were sampled in elliptical annuli. Since the resolution with which the images were scanned by the APM machine was roughly $1 \text{ pixel arcsec}^{-1}$, an image of diameter 1 arcmin would have a semimajor axis of 30 pixel. We therefore decided to sample the galaxies on 30 elliptical annuli, all of which share the position angle and ellipticity which were derived for the whole image. This ensures independence of inclination, although obviously there is less information in the edge-on images. For images whose diameter is larger than 1 arcmin the annuli are spread out over the entire image, with a constant spacing in the major axis. The number of sampling points on each annulus starts at 32 for the innermost and rises through 64, 128 and 256 up to 512 for the outermost. This constitutes some oversampling in most annuli.

All subsequent feature extraction is performed on the sampled annuli. The full image as such is no longer used. In the bottom half of Fig. 3 we show two images of the same galaxy – one is the sky-subtracted and star-removed image, while the other is its reconstruction from the sampled annuli. The reconstructed image is rotated (since the sampling begins always on the automatically determined major axis), and the central few pixels are missing since they were not sampled. The original image diameter was much larger than 1 arcmin , so we undersampled it considerably with the 30 ellipses. For this reason the reconstructed image has a significantly lower resolution than the original image. Nevertheless, it is undoubtedly the same galaxy, and, as we show below, the loss of information in such cases does not prevent the ANN from functioning well.

5 FEATURE EXTRACTION

Throughout this paper our approach to the classification problem is trying to mimic the human expert. We therefore needed to extract global as well as detailed parameters, according to our understanding of what is morphologically significant. We have developed considerable software for this task and all the features extracted for the ANN are prepared in a fully automated manner. In the description that follows we make repeated references to four representative galaxies, whose images are shown in Fig. 4. We begin by describing the parameters and how they were extracted from the images, and then proceed to analyse their correlations with type.

5.1 Global parameters

The first two global parameters we have are the apparent ellipticity and the surface brightness of the image, which are calculated by the reduction program. The other source of global parameters is the light profile, which is a plot of average annulus intensity versus annulus number (in the range 1–30). The light profiles of the selected four galaxies

are shown in Fig. 5. Although the appearance to the eye of the two edge-on galaxies is quite different from that of the face-on galaxies, the similarities between the light profiles of the two early-type galaxies (f643g004 and f646g057) are clear. The profile of f646g004 is quite unique, as the internal ring is so prominent that the light profile actually *rises* around annuli 7–10.

The general quality of the light profiles is too poor for a fitting procedure of either an $r^{1/4}$ law or an exponential disc to be carried out, and so we have no means of determining the bulge-to-disc ratio, which is a very important indicator of type (Simien & de Vaucouleurs 1986). We therefore turn to other parameters derived from the light profile. A concise way to describe the light profile is to measure concentration indices (cf. Doi, Fukugita & Okamura 1993). We used the definition

$$C_\alpha = \frac{\int_0^{30\alpha} r I(r) dr}{\int_0^{30} r I(r) dr}, \quad (3)$$

where r is the radius in normalized units (1–30), $I(r)$ is the average intensity of the corresponding annulus, and α is a fraction. We calculated nine such indices for α values of 0.1, 0.2, ..., 0.9, and added as a tenth parameter the ratio of the radius enclosing 75 per cent light to the radius enclosing 25 per cent light, as suggested by de Vaucouleurs (1963).

While the concentration indices will correctly convey the similarity between the two early-type galaxies, they will not tell us much about the affinity of the two spirals (f646g004 and f635g008). We therefore need other parameters to tell us that. Although the differences between their light profiles are considerable, the two spirals do have in common a certain feature: both of their light profiles drop fast at low annuli, whereas those of the early-type galaxies are nearly flat in the first few annuli before beginning to decrease significantly. This is a reflection of the size of the bulge, which in many galaxies in our sample is burnt out. Since we cannot measure the bulge-to-disc ratio directly, the alternative we are left with is measuring the size (i.e. number of annuli) of the ‘flat’ region in the light profile. The definition we adopted for the size of the flat region is as follows. We follow the slope of the light profile and find its steepest value. We then define 25 per cent of this value to be the threshold, and the radius at which the slope first exceeds this threshold is taken to be the size of the bulge. If the light profile starts with a sharp rise, we take the bulge size to be zero.

Two other parameters may be of importance. One is the shape of the outer part of the profile, which is much straighter for the spirals than for the early-type galaxies. A way of quantifying this observation is to fit straight lines to the middle and outer third of the profile and look at the ratio of the slopes. Another parameter of interest, suggested by D. Lynden-Bell, comes from looking at a plot of $r \times I(r)$ versus r , rather than the original profile [$I(r)$ versus r]. This kind of plot for the four galaxies is shown in Fig. 6. It can be seen that the peak of this curve occurs at values of r that increase with type, and the two early-type galaxies have almost the same peak position. Here again the prominent ring of f646g004 plays a part in shifting the peak of its curve to the left. We take the position of the peak as another parameter. The total number of global parameters is therefore 15 [ellipticity, surface brightness, 10 concentration indices and




Figure 4. Images of four galaxies representing early types (left) versus late types (right), and edge-on (bottom) versus face-on (top) inclinations.

'bulge size', slope-ratio and the peak position of $r \times I(r)$ versus r].

5.2 Detailed parameters

All detailed parameters should be available from the sampled annuli, and in Fig. 7 we show the annuli of each of the four galaxies, with r going up along the positive y -axis. Detailed parameters are by nature available only for face-on galaxies, since the amount of information drops sharply as the inclination increases. The plot for the edge-on early-type galaxy is completely dominated by its bulge, and it actually appears to have small-scale structure comparable to that of the edge-on spiral. Clearly, only the plots of the face-on galaxies can be trusted to supply us with useful information for classification.

We therefore defined a quantitative criterion for the distinction between edge-on and face-on galaxies. Inclination is not the only factor that plays a part here. Large images can

have high inclinations but still contain much detail, whereas small images already contain little information at moderately high inclinations. We therefore decided to look at the semiminor axis. If the semiminor axis is too small, many of the annuli will overlap as they pass near the bulge and we will get certain features smeared over many annuli, although their actual sizes are quite small. For this reason our criterion for 'edge-on image' is that its semiminor axis is less than 15 pixel long (i.e. the maximum number of annuli allowed to overlap on any given pixel is 2).

Splitting our sample in view of this criterion means that we can try to work out detailed parameters only for less than 70 per cent of the sample. For this part of the sample we found two promising sources of parameters. First, there is the parameter used by Schweizer (1976), who investigated six giant spirals in great detail and defined an *arms-to-disc* ratio in the attempt to quantify the ratio of arm-strength to the underlying disc. Following the definition in Mihalas & Binney (1981), we looked at each elliptical annulus and took




Figure 5. Light profiles of the four selected galaxies.

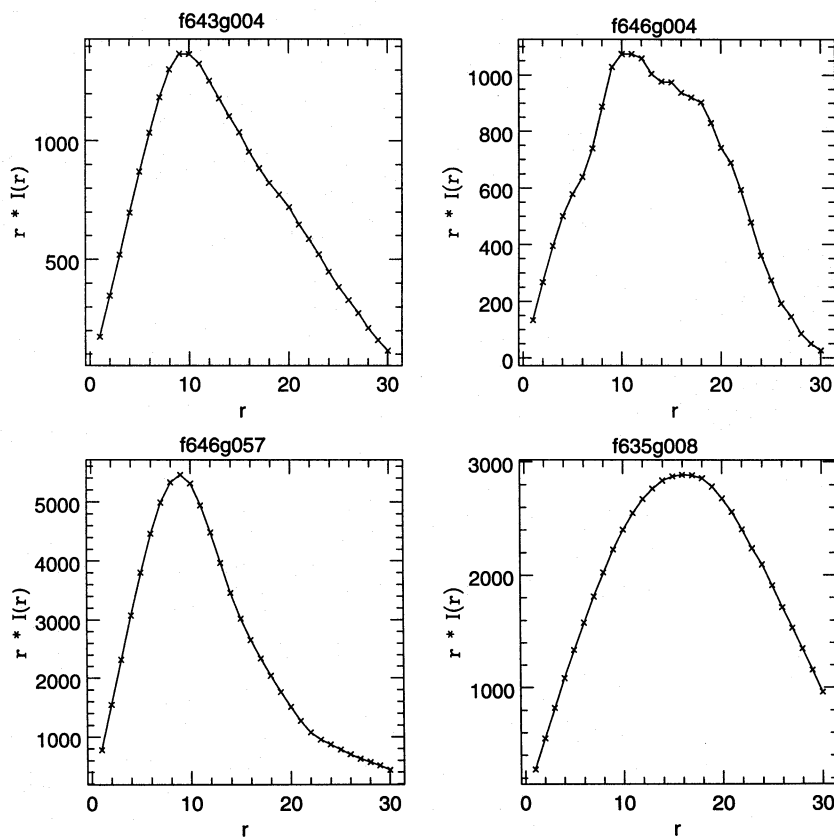


Figure 6. $r \times I(r)$ versus r curves. The curves of early-type galaxies peak at lower r values.

as the disc component (I_d) the average of the two lowest intensity minima lying at least 90° away from each other. The average of the annulus (I_a) was then calculated, and the arms-to-disc ratio was defined as

$$AD = \frac{I_a - I_d}{I_d}. \quad (4)$$

This ratio was calculated for each annulus separately, and then we made six combinations of the 30 ratios: their mean over all annuli, and the means over five groups of successive annuli (3–7, 8–12, 13–17, 18–22 and 23–27).

The other source of detailed parameters is the spiral arms. Hubble (1926, 1936) used them as one of the major parameters in classification, yet gave qualitative rather than quantitative definitions. Quantifying the arms implies tracing them in the first instance. We developed a computer code for doing just that. Our starting point is a projection in polar coordinates of all annuli (each annulus stretched/shrunk to 128 points in the azimuthal direction). We then define as the 'local neighbourhood' of any given pixel the eight pixels nearest to it (one on either side of it in the same annulus and the nearest three in either adjacent annulus). Next, we cross out from this projection all pixels with more than two neighbouring pixels of higher intensity. This gives us a basic trace of bright patches of arms, but the arms the human eye 'sees' look at this stage quite fragmented. In Fig. 8 we show a trace-map of the arms of the face-on spiral we discussed earlier (f646g004), as picked up by our software. The arms stand

out, but clearly some of them are fragmented, and there is noise. We employ a connecting algorithm to link arm fragments, subject to criteria of distance and orientation. Finally, we measure the average length of the resulting arms, the number of arms and their intensity, which is measured as the average ratio of intensity of pixels belonging to the arm to the average intensity of their respective annuli. The total number of detailed parameters we end up with is nine (six arm-to-disc ratios and three arm parameters). For the edge-on images we kept all detailed parameters fixed at a value of zero. Table 3 summarizes our parameters and their definitions.

5.3 Correlations of parameters with types

We checked to see whether there are good correlations of our various parameters with the Revised Hubble type. In Figs 9–12 we show plots of some parameters versus the corrected mean type. Figs 9 and 10 describe global parameters for all galaxies. Figs 11 and 12 describe detailed parameters for face-on galaxies only. The distribution of values is significantly different from Gaussian for many types, and so for each type we calculated the median value and drew a vertical bar from the value representing 32 per cent to the value representing 68 per cent of the galaxies in that type. The parameters chosen appear to be correlated with type, but there is considerable scatter and many of the correlations are non-linear. This is an indication that our choice of parameters requires a non-linear automated classifier, such as an ANN.

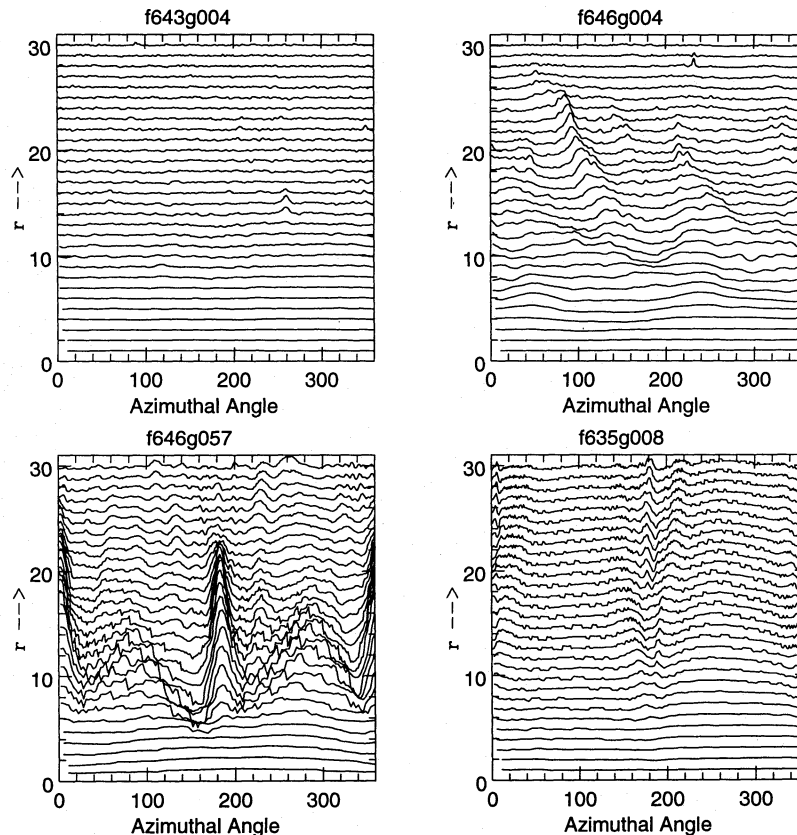


Figure 7. All annuli of the four selected galaxies, artificially shifted so that the innermost annulus is at the bottom of each plot.




Figure 8. f646g004: a trace of the spiral arms. Compare with the original and reconstructed images in the lower half of Fig. 3.

Table 3. Definitions of morphological parameters for the ANN.

Number	Parameter Name	Definition
1	Ellipticity	$1 - b/a$, as measured by the reduction program
2	Surface Brightness	average pixel intensity over the whole image
3	Bulge Size	radius at which light profile slope first exceeds threshold
4	Slope Ratio	ratio of slopes of straight lines fitted to the middle and outer 1/3 of the light profile
5	Peak rI	radius where peak of $rI(r)$ vs. r occurs
6	C_{10}	concentration index, $\alpha = 0.1$ (see eq. 3)
7	C_{20}	concentration index, $\alpha = 0.2$
8	C_{30}	concentration index, $\alpha = 0.3$
9	C_{40}	concentration index, $\alpha = 0.4$
10	C_{50}	concentration index, $\alpha = 0.5$
11	C_{60}	concentration index, $\alpha = 0.6$
12	C_{70}	concentration index, $\alpha = 0.7$
13	C_{80}	concentration index, $\alpha = 0.8$
14	C_{90}	concentration index, $\alpha = 0.9$
15	$C_{75/25}$	ratio of radius enclosing 75% of total intensity to radius enclosing 25% of it
16	AD_{all}	arms-to-disk ratio, averaged over all annuli
17	AD_{3-7}	arms-to-disk ratio, averaged over annuli 3 – 7
18	AD_{8-12}	arms-to-disk ratio, averaged over annuli 8 – 12
19	AD_{13-17}	arms-to-disk ratio, averaged over annuli 13 – 17
20	AD_{18-22}	arms-to-disk ratio, averaged over annuli 18 – 22
21	AD_{23-27}	arms-to-disk ratio, averaged over annuli 23 – 27
22	Arm Number	number of arm segments found by the arm-seeking algorithm
23	Arm Length	average length of arm segments
24	Arm Intensity	average relative intensity of arm segments

Of course, not every parameter we use needs to be well correlated with type. For example, high ellipticity can tell us that a certain image is not that of an elliptical galaxy, and yet there is no clear correlation between ellipticity and type in general. There are some curiosities too. The ‘bulge size’ seems to go down rather than up as we go to types earlier

than type 0 (S_0/a). This may be due to some artefact in our sample, or it maybe reflects the difference between our definition of ‘bulge size’ and the accepted definition which relies on the fit to $r^{1/4}$ profiles. Type 10 (Im) galaxies do not fit very well with the trends set up by other types, for most parameters. This may be due to the very small number of




Figure 9. Correlations of light-profile parameters with corrected mean type.




Figure 10. Correlations of light-concentration indices with corrected mean type.

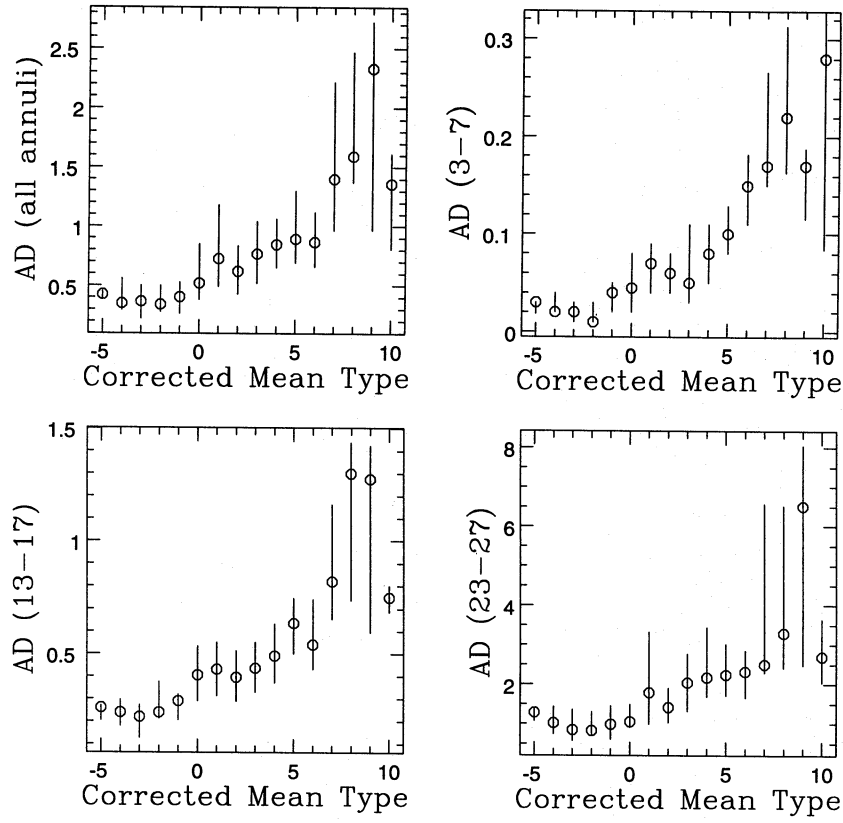


Figure 11. Correlations with arms/disc parameters with corrected mean type.




Figure 12. Correlations of spiral arm parameters with corrected mean type.

type 10 galaxies in our sample and to the fact that some of them are severely overexposed. There is noise too in the arm parameters, especially in the average arm-length parameter, but the overall correlation is good enough for us to include it. In the arm parameters we have non-zero values for the early-type galaxies, which are the result of noise and superposed stars (which were not removed by the reduction program) being interpreted by the software as arm fragments. There are, of course, internal correlations between our parameters, and these were dealt with as discussed below.

6 THE ANN CLASSIFICATION PROCEDURE

6.1 Combining parameters to form input parameter sets

The feature extraction stage produced 24 parameters, which we defined as our 'parameter set I'. However, we had good reasons to expect equally good results if we use less parameters. There are obvious correlations among the parameters we chose. Clearly, we do not need all 10 concentration indices or all six arms-to-disc ratios. In addition, the ANN performance could be degraded if we use too many inputs, as this means the number of free parameters ('weights') it is using is very large. This increases the chance of the ANN settling into a local (rather than the global) minimum (the considerations for constructing the ANN will be elaborated upon below).

In order to examine possible correlations in the data, we used principal component analysis (PCA). This is a linear technique whereby a parameter space of some dimensionality is rotated to its principal axes, defined as the orthogonal axes along which the variances are maximal. It is usually possible to represent a high percentage of the variance in the data by a number of principal components that is much smaller than the original dimensionality of the data (see Lahav et al. 1995, in preparation, for more details). We employed this technique separately on each of the four 'families' of input parameters, namely the light-profile parameters (bulge size, slope ratio and peak r/I), the concentration indices, the AD ratios and the spiral-arm parameters. For the purpose of the PCA, each parameter was normalized to have zero mean and unit variance, so as to avoid one large-valued parameter overwhelming the rest and biasing the intercorrelations between the parameters. As a result of this analysis we chose the following combinations of parameters to make up our compressed 'parameter set II': two linear combinations of the three light-profile parameters (covering 93 per cent of the variance); three combinations of the 10 concentration indices (98 per cent variance); three combinations of the AD ratios (93 per cent variance) and all three arm parameters. We added the ellipticity and the surface brightness without change, so 'parameter set II' contains 13 parameters.

In view of the possible application of our methods to small images of distant galaxies, we wanted to know the effect of

removing the detailed parameters and leaving only those parameters which can be measured for small, distant images. We therefore defined a set of parameters which included only the global parameters: 'parameter set III' includes the ellipticity, the surface brightness, the two PCA-compressed light-profile parameters, and the three PCA-compressed concentration indices – seven parameters in all. Table 4 summarizes our various parameter sets.

All the parameters in all of the sets were then rescaled so as to fit into the range $[0, 1]$ by the transformation

$$X_{\text{new}}^j = \frac{X_{\text{orig}}^j - X_{\text{min}}^j}{X_{\text{max}}^j - X_{\text{min}}^j}, \quad (5)$$

where the superscript j denotes which parameter is being rescaled; the 'new' and 'orig' subscripts refer to the rescaled and original values of that parameter, respectively; and the 'min' and 'max' subscripts refer to the minimal and maximal values (respectively) of that parameter over the entire sample.

6.2 ANN configurations

6.2.1 Designing the ANN architecture

The architecture of the ANN is the full description of the way different nodes in the ANN are connected to each other. The number of nodes in the input layer is dictated by the number of parameters in the chosen training set. There was a single output node in all runs. We used sigmoids for the transfer functions from one layer to the next, defined as follows:

$$f(\mathbf{x}, \mathbf{w}) = \frac{1}{1 + \exp(-\mathbf{w} \cdot \mathbf{x})}, \quad (6)$$

where \mathbf{x} is the vector of data entering the current layer, and \mathbf{w} is the weights vector. The range of the sigmoid function is $[0, 1]$, so the types were transformed from their original values (we used only types in the range $[-5, 10]$) to the range $[0, 1]$ in a linear manner.

The number of hidden-layer nodes is not specified from the outset and is changed to get the best overall results. For most runs we used parameter set II, which had 13 inputs, so we decided to try one architecture with 13 hidden nodes and one with five, in order to see if such a significant difference in the number of weights would change the results.

The ANN is a minimization scheme in the multidimensional space of the weights (the free parameters of the ANN). The value it minimizes is the rms difference between the desired output (e.g. an expert type) and the actual ANN output:

$$E = \frac{1}{2} \sum_i (D^i - O^i)^2, \quad (7)$$

Table 4. Parameter sets for training the ANN.

Set Number	No. of Parameters	Contents
I	24	A full set including all parameters
II	13	PCA-compressed version of set I
III	7	Global parameters only, PCA-compressed

where D^i is the desired output for galaxy i , and O^i is the actual ANN type for that galaxy. The summation is over all galaxies in the training set. There is an interplay between the degree of complexity, allowing a more refined classification as the number of weights increases, and the existence of local minima in the error. Thus the chance of the ANN getting stuck in them increases with the number of weights. For this reason we use a regularization term ('weight decay'). This term is included by changing the error we are minimizing to

$$E_{\text{tot}} = E + \frac{\epsilon}{2} \sum_j w_j^2, \quad (8)$$

where E is defined in (7), and the summation is over all weights in the ANN. We fixed the value for ϵ by following the recipe given in Lahav et al. (1995, in preparation), which is based on Mackay (1992). We found optimal values of $\epsilon = 0.003$ for ANNs with five hidden nodes, and $\epsilon = 0.01$ for ANNs with 13 hidden nodes, from a set of 'design' runs. These were carried out over a grid of ϵ values in the range $[0.0, 1.0]$, for each of the two architectures we used and using parameter set II with the corrected mean types of all six experts. The optimal values we found for ϵ also gave the smallest rms difference between the ANN and the desired outputs.

6.2.2 The classification runs

On top of the various type sets we chose for each of the architectures, the process of training the ANN requires splitting the parameter set in two to form a training set and a test set. Each of these sets contains, for every galaxy, the values of all the input parameters and the desired value(s) of the output(s). The ANN 'learns' from the training set, trying to get as close as possible to the desired output(s) for each galaxy. When the training phase is over, we can use the resulting weights to classify data, by presenting an input vector to the ANN and observing the output. ANNs with a single (so-called 'analogue') output produce a continuum of output values, which simply give us a predicted T -value for each galaxy. On the other hand, one may use a multiple-output architecture which allows one to get classifications in terms of Bayesian a posteriori probabilities of belonging to any given type (see Lahav et al. 1995, in preparation). However, galaxies form a sequence and we wanted to avoid binning them artificially into a certain number of classes. For the purposes of this paper we decided not to make use of multiple output ANNs at all.

In planning the various configurations of the ANN we had to solve several problems.

(i) We had a relatively small number of galaxies. Our sample contained 830 galaxies, which seems to be quite a large number. However, there were 16 different types, so some types were not well-represented. Moreover, we had to split our sample into a training set and a test set. Due to this severe shortage of training patterns we split the data randomly in four and made four pairs of training-test sets, each time picking a different combination of three quarters to make up the training set and the remaining quarter as the test set. As a result we had to run each ANN four times for

each configuration just to get a full classification of all galaxies.

(ii) We had 11 different sets of desired outputs (type sets), as explained above. For each of the sets 4–9 there was only one source of types (the classifications of the chosen expert). However, sets 1–3 were made from means of several experts, and we felt we were losing information by training the ANN on the mean type only. We approached the problem by enlarging our training set artificially, including in it each galaxy not once but rather as many times as the number of classifications it received by different experts. All of its appearance in the enlarged set had exactly the same input values, and the only difference between various entries of the same galaxy was the desired type, supplied by a different expert each time. Thus we included all the information supplied by the experts and also gave more appearances to galaxies which more experts classified. In the test set, however, each galaxy had to have a single appearance, with one type (the 'true answer'). We enlarged our training sets this way when constructing them in conjunction with types sets 1, 2 and 3. In Table 6, which summarizes the results of various runs, the 'training set size' column is the approximate number of patterns in a typical training set. This number is not unique, as we composed the training sets for each run from three quarters which were only approximately of the same size and so the numbers of patterns quoted in the table is just an average. A possible outcome of the fact that some galaxies appear more often than others is a change in the frequency of types in the enlarged training set, as compared with the original set. We found that this effect is marginal in our sample. The runs using the enlarged training sets are numbered 3, 4, 7, 8 and 23–26 in Table 6. As a consistency check we also made runs with a training set that was not enlarged, using type set 2. These were numbered 5 and 6.

(iii) The ANN starts running from a random set of weights. Even with the regularization of the weight decay term, it is conceivable that some initial weights will lead the ANN to minimum values of the error that are not quite the global minimum it seeks. To get a reliable answer one has to run the same ANN several times, each time initializing the weights at a different set of values, and compare the results. We used 10 different randomizations for each of the four combinations of quarters, and so each result we quote for a given configuration is the outcome of 40 runs of that configuration. We used the same 10 randomization seeds for all configurations in order to have a uniform basis for comparing different runs.

(iv) We needed a criterion for stopping the training process. The error on the training set declines monotonically due to the ANN achieving better and better fits to the training patterns. The error on the test set, on the other hand, is the indicator of the ANN's ability to correctly classify data it has not been trained on, i.e. to be a good classifier for the vast number of galaxies which were not classified by human experts. For this reason we ran the ANN in all configurations for 300 iterations, and the set of weights giving the minimal error over the *test set* was kept and later used for classifying.

The list of configurations we ran appears in Table 5. The architecture is specified as $N_{\text{inp}}:N_{\text{hid}}:N_{\text{out}}$, where N_{inp} is the number of inputs, N_{hid} is the number of hidden nodes, and N_{out} is the number of outputs. In two particular configura-

tions we used no hidden units at all (denoted by a value of 0 for N_{hid}). The code used was a quasi-Newton (see, e.g., Press et al. 1993) ANN package written by B. D. Ripley of Oxford University, with some adaptations. In the case of no hidden units the ANN had either purely linear or purely sigmoidal connections. Typically, each run took between 1 and 10 CPU minutes to converge on a SUN Sparc 10 workstation. Once the ANN finished training, the actual classification of fresh data took about 1 CPU second for the whole sample.

Table 5. ANN configurations.

ANN Architecture	Parameter Set	Types Set
24 : 5 : 1	I	2
24 : 13 : 1	I	2
13 : 0 : 1	II	2
13 : 5 : 1	II	1,2,3,4,5,6,7,8,9
13 : 13 : 1	II	1,2,3,4,5,6,7,8,9
7 : 3 : 1	III	2
7 : 7 : 1	III	2

6.3 Results

6.3.1 General results and rms dispersions

In order to compare the results of the ANN with the results we found for the human experts in Paper I, we use here the same measure of classification success that we used there, namely the rms dispersion, taken in this paper between ANN output and the desired output (the type of the galaxy as dictated from outside). Table 6 summarizes the results for all the configurations we tried. It specifies for each configuration the rms dispersion between the ANN and the relevant type set, and the average over all galaxies of the rms dispersion between different runs of the ANN (the ‘internal scatter’ of the ANN). The rms dispersion is defined as

$$\sigma^2 = \frac{1}{N_{\text{gal}}} \sum_{i=1}^{N_{\text{gal}}} (T^i - T_{\text{ann}}^i)^2, \quad (9)$$

where T is the relevant expert type, in the range $[-5, 10]$, and T_{ann} is the ANN type, which is expanded from the range $[0, 1]$ to the range $[-5, 10]$ using a linear transformation. The

Table 6. Results of runs as measured by the rms dispersion between the ANN type (averaged over 10 runs) and the expert types. CMT stands for corrected mean type, and SMT for straight mean type. In runs 3, 4, 7, 8, 23, 24, 25 and 26 use has been made of the enlarged training sets, as described in the text. Parameter sets are as in Table 4.

Run Number	ANN Architecture	Parameter Set	Types Set	RMS Dispersion	Internal Scatter	Training Set Size
1	13 : 0 : 1 Lin.	II	2 (CMT of all experts)	2.04	0.07	620
2	13 : 0 : 1 Sig.	II	2 (CMT of all experts)	2.00	0.06	620
3	13 : 5 : 1	II	2 (CMT of all experts)	1.83	0.35	3200
4	13 : 13 : 1	II	2 (CMT of all experts)	1.84	0.28	3200
5	13 : 5 : 1	II	2 (CMT of all experts)	1.85	0.25	620
6	13 : 13 : 1	II	2 (CMT of all experts)	1.90	0.11	620
7	13 : 5 : 1	II	1 (SMT of all experts)	1.79	0.35	3200
8	13 : 13 : 1	II	1 (SMT of all experts)	1.81	0.28	3200
9	13 : 5 : 1	II	3 (CMT of RB,HC,GV)	1.92	0.28	620
10	13 : 13 : 1	II	3 (CMT of RB,HC,GV)	1.94	0.22	620
11	13 : 5 : 1	II	4 (RB)	1.94	0.23	570
12	13 : 13 : 1	II	4 (RB)	1.96	0.12	570
13	13 : 5 : 1	II	5 (HC)	1.96	0.27	610
14	13 : 13 : 1	II	5 (HC)	2.00	0.13	610
15	13 : 5 : 1	II	6 (GV)	2.24	0.33	350
16	13 : 13 : 1	II	6 (GV)	2.30	0.12	350
17	13 : 5 : 1	II	7 (AD)	1.89	0.27	610
18	13 : 13 : 1	II	7 (AD)	1.95	0.11	610
19	13 : 5 : 1	II	8 (JH)	2.31	0.32	620
20	13 : 13 : 1	II	8 (JH)	2.34	0.17	620
21	13 : 5 : 1	II	9 (vdB)	2.18	0.31	410
22	13 : 13 : 1	II	9 (vdB)	2.25	0.13	410
23	7 : 3 : 1	III	2 (CMT of all experts)	2.22	0.45	3200
24	7 : 7 : 1	III	2 (CMT of all experts)	2.15	0.40	3200
25	24 : 5 : 1	I	2 (CMT of all experts)	1.78	0.39	3200
26	24 : 13 : 1	I	2 (CMT of all experts)	1.81	0.28	3200

summation is over all galaxies, and the superscript i denotes a particular galaxy.

The following comments can be made about this table.

(i) For the purely linear ANN with no hidden units (run 1) one may write down an analytical solution (Lahav et al. 1995, in preparation). Although the ANN in this case is much simpler and faster to converge, it does significantly worse than its non-linear counterparts. This supports the view that our classification problem is not linear, and that the use of classifiers such as ANNs is well justified. The use of a hidden layer contributes a lot, as can be seen from the fact that a non-linear ANN, when run without hidden nodes (run 2), gives poorer results than those configurations with hidden nodes.

(ii) While it is obviously important to include hidden nodes in the ANN architecture, their exact number has little effect on the results. This is so, because we adapt the weight decay coefficient ϵ to the architecture. This means we can save a lot of computer time by choosing a relatively small number of hidden units (say, a third of the number of inputs) when constructing the ANN.

(iii) The ANN fits the mean of all experts (either straight or corrected mean) better than it does any of the experts taken separately, at least when its configuration includes some hidden nodes. It could be that in calculating the mean type the personal tendencies of each expert are counteracted by the other experts to form a more balanced set of classifications. It is important to remember that it is the ability of the ANN to replicate classifications by a human expert that is measured here, not the reverse. A larger rms dispersion between the ANN and an expert means that given the parameters, the ANN was less successful in learning that expert's classification scheme.

(iv) Comparing to Paper I, where we found the overall rms dispersion of the experts to be about 1.8 types, the ANN achieves a comparable rms dispersion, at least when measured against the mean types of all experts (runs 3–8). The runs utilizing the enlarged training sets (runs 3, 4, 7 and 8) gave slightly better results than those that used the original training set, but the difference is too small to be considered significant.

(v) The agreement of the ANN with the averaged classifications of RB, HC and GV (runs 9 and 10) is slightly worse than the agreement with the corrected means of all six experts, but better than the results for the sets utilizing classifications of any one of them separately (runs 11–16). This may again reflect the ‘mitigating’ influence of taking the mean type of several experts. Note also that these three experts had the tightest correlations between their classifications (see Paper I).

(vi) Use of all 24 parameters (runs 25 and 26) has no advantage over using the PCA-compressed set of 13 parameters, in terms of successful classifications. On the other hand, use of all 24 parameters forces us to utilize more weights in the ANN architecture which only results in much longer computer runs, so our natural choice is to concentrate on the set of 13 parameters.

6.3.2 Interpreting the ANN classifications

Analysis of the ANN classifications in terms of a single number (the rms dispersion) does not tell the whole story.

For a more detailed study of the way the ANN worked, we looked closely at run 3 (13:5:1 architecture, corrected mean types). In Fig. 13 we show four plots. The bottom-left plot shows the frequency of corrected mean types in our sample of galaxies. The bottom-right plot shows the frequency of types given by the ANN in our sample. It is clear that the ANN avoided both ends of the scale, smoothed the overall distribution and was most inclined to classify galaxies into type 4, which is the most frequent in the sample. The top-left plot describes the rms dispersion between the ANN and the corrected mean type *for each type separately*. It is clear that the minimum of the rms plot corresponds to the most frequent types (3–5) in the sample. The erratic behaviour of this plot at the early-type end is due to the fact that most of these galaxies were assigned either type –5 or type –2 by the experts, and the resulting mean type was somewhere between these two extremes. This meant that for these types the ANN had to force its output to be what no input pattern actually demanded it to be. The ANN could not always quite do this transition, and while for some cases its designated class matched the mean well, in others it did not do so well. The top-right plot shows how the rms dispersion between the ANN and the experts changes with ANN type. This can be taken to be the measure of reliability of the ANN classifications, and it is rather constant, with both ends of the scale artificially dropping to zero, since the ANN did not classify any galaxy as either types –5, 9 or 10. What the plots in Fig. 13 show clearly is that the ANN is inclined to take ‘safe bets’ – in general, the more patterns of a given type it sees, the more likely it is to choose that type when in doubt. To examine further the ANN performance, we show in the left panel of Fig. 14 a scatter plot of ANN type versus corrected mean type for all galaxies in our sample. The error bars indicate the standard deviation of the classification of each galaxy by the ANN over 10 runs (the internal scatter of the ANN). The diagonal lines are equal-error lines. The tight correlation around types 3–5 was expected from the previous figure, but in this figure, we also see what happens at the early and late ends. Whereas the ANN classification still roughly follows – on average – the mean expert type for early mean types, it fails almost completely in classifying the very late-type galaxies ($T > 7$). In the right-hand part of Fig. 14 we show a histogram of the rms dispersion between the ANN types and the corrected mean types, as a function of internal scatter in the ANN classifications. There is a correlation between these two quantities, although it is rather noisy. This can mean that *like a human expert* the ANN tends to make mistakes when it is less sure what type to attach to a certain galaxy. It is worth mentioning that while a human expert can exercise the right not to classify a certain galaxy, the ANN has no such freedom and is forced to classify all galaxies.

Another interesting question is whether there is any limiting ‘hard-core’ error value below which the ANN cannot go. To answer this, we plotted in Fig. 15 the cumulative percentage of galaxies for which the error in classification is not more than a given value, against the error. The curve is very smooth and shows clearly that there are errors of all values. The curve is concave, so that as the error increases the marginal growth in the number of galaxies decreases. Indeed, 54 per cent of all galaxies have errors of up to one type, and 80 per cent have errors of up to two types. These values are




Figure 13. Top: average rms dispersions of the ANN classifications versus corrected mean type (left) and ANN type (right). Bottom: frequency of corrected mean types (left) and ANN types (right) in our sample.




Figure 14. Left: ANN classifications versus corrected mean type. Error bars indicate scatter in ANN classification over 10 runs. Diagonal lines correspond to differences of 0, ± 1 , ± 2 and ± 3 types between the ANN and the corrected mean type. Right: histogram of rms dispersion versus ANN internal scatter.

comparable to the values we calculated for the human experts in Paper I.

6.3.3 Dependence on image diameter and inclination

Fig. 16 shows the dependence of the rms dispersion between the ANN type and the corrected mean type on image diameter and axial ratio. There is a trend towards smaller

rms dispersion as the size of the image increases, but it is very weak and noisy. This seems to indicate that our sampling technique gives us a comparable quality of features from all image sizes in the sample. On the other hand, the axial ratio dependence seems to suggest that the ANN is doing much better on *edge-on* images, which is rather surprising. This, however, can be explained by the fact that most edge-on images are of types 3–5, where in general the classification




Figure 15. Cumulative plot of ANN classifications correct to within a given error versus classification error.




Figure 16. Histograms of rms dispersions versus image diameter (left) and axial ratio (right).

errors are low, while many face-on images belong to early-type galaxies, for which we have a higher rms dispersion. At any rate, the trend one would expect – that of rms dispersion increasing with image inclination – is not there. Again, this means that our sampling is good in that it extracts reliable features from all images, regardless of their apparent ellipticity.

6.3.4 Analysis of severe misclassifications

We looked at a subset of the sample for which the ANN missed the corrected mean type by at least three types

(roughly a 2σ error). There were 74 such galaxies (9 per cent of the sample), and in the top part of Fig. 17 we show their distribution versus corrected mean type. While the absolute number of misclassifications is relatively stable across all types (with some increase towards both ends), the *fraction* of misclassifications is evidently much higher towards both ends of the scale, especially for very late types where it shoots up for types 9 and 10. The problem of classifying at both ends of the scale may be due to our choice of parameters (failing to describe those types of galaxies properly), or to small-number statistics (see Fig. 13 for class frequencies). We think it is due to both, and in order to check this we




Figure 17. Frequencies of badly misclassified galaxies (top) and of the best classified galaxies (bottom) versus corrected mean type, as frequencies out of each type separately. Compare with class frequencies in Fig. 13.

again show, in Figs 18–21, the distributions of several parameters we used versus corrected mean type, but this time we add each of the 74 badly misclassified galaxies as single points on these plots. The bulge sizes for most of these galaxies are far outside the range of values typical for their respective types, and there are also large discrepancies in the slopes ratios and peak r/l . These problems emerge again in the light concentration indices, where many of these galaxies are out of the scale of any type whatsoever. This means that globally something is wrong with these images – e.g. their sizes were overestimated, leading to a wrong light profile that includes a lot of sky instead of only the galaxy, or they were overexposed, leading to exceedingly large bulges. It is difficult to draw conclusions from the detailed parameter plots (Figs 20 and 21), since some of these galaxies were edge-on and consequently their detailed parameters were set to zero to begin with. Nevertheless, it seems that some early-type galaxies had ‘arms’ (apparently noise in the image), while the very late-type galaxies were underestimated for arms and arm-to-disc ratios. We can try to settle the question of which type of parameters was responsible for these misclassifications by plotting the fraction of misclassified galaxies against their apparent ellipticity. Since the detailed parameters were unavailable for edge-on images, this kind of plot should tell us whether their absence made the edge-on images more susceptible to large errors in classification. In Fig. 22 we show two superposed histograms: the solid line denotes frequencies of axial ratios in the entire sample (expressed as a percentage of the total sample), while the dashed line

denotes the badly misclassified galaxies (as a percentage of these 74 galaxies). The two plots are very similar, and this is taken to imply that the lack of detailed parameters did not cause more misclassifications – on the contrary, for very low axial ratios (0.2 and less) there are less badly misclassified galaxies than there are such galaxies in the whole sample.

This leaves us with the global parameters as the major culprits. We have looked at the images of these 74 galaxies one by one and tried to identify the problems. For about 30 of them there was not one obvious source for the large error. However, for the others we found that 23 were overexposed, which led to overestimation of bulge sizes and influenced the entire light profile; in 10 early-type images the sky was not cut at the correct threshold, leading to an underestimation of the bulge and causing spurious structure in the external regions to be regarded as arms; there were three interacting pairs and six or seven peculiarly shaped galaxies, which our standard parameters were not designed to describe; and there was one case where the image quality was very poor and the resulting parameters were degraded. The implication is that when applying this kind of classifier to new data one has to take care and monitor the images by eye as a safety check, or alternatively devise some automated means for detecting the problems we encountered.

6.3.5 Analysis of the best classifications

We have defined the ‘best classified galaxies’ as the collection of all galaxies whose ANN type differs by less than 0.5 types




Figure 18. Light profile parameters: badly misclassified galaxies (crosses) and the dispersion of parameter values in the whole sample, versus corrected mean type (circles).




Figure 19. Light concentration indices: badly misclassified galaxies (crosses) and the dispersion of parameter values in the whole sample, versus corrected mean type (circles).




Figure 20. Arms-to-disc ratios: badly misclassified galaxies (crosses) and the dispersion of parameter values in the whole sample, versus corrected mean type (circles).




Figure 21. Arms parameters: badly misclassified galaxies (crosses) and the dispersion of parameter values in the whole sample, versus corrected mean type (circles).




Figure 22. Distribution of axial ratios in the whole sample (solid line) and in the subset of badly misclassified galaxies (dashed line).

from the expert type. 231 galaxies (28 per cent of the sample) met this criterion, and the lower half of Fig. 17 shows their distribution with type. This distribution closely follows the distribution of types in the entire sample and, as expected, the most frequent types in the sample are those most represented in the 'best classified' subset. Note that no galaxies of types -5 , 9 and 10 belong to this subset.

7 DISCUSSION

We describe an automated process starting with the digitized images of galaxies and ending with their morphological classification. The sample was selected from the APM Equatorial Catalogue (Section 2), but the techniques described in this paper can be easily applied to other sources of galaxy images. Classifications for this sample were provided through collaborative efforts with six experts (Section 3), and can be used to classify automatically many other galaxies. The data reduction (Section 4) and feature extraction (Section 5) were done by our software. The ANN was trained on various sets of types and performed very well. The rms dispersion between the ANN types and the corrected mean type is comparable to the overall rms dispersion between pairs of experts (as found in Paper I). Analysis of the way the ANN classifies (Section 6) showed that it is more or less reliable to the same degree for most types, but does not classify any galaxy into types -5 , 9 or 10 . Its error compared to the mean expert type is smallest for the type range $3-5$, which is the most frequent in the sample. These can be taken as measures of the ANN reliability in future classifications of many more images, predominantly of as yet unclassified galaxies. Another extension of this work can be the application of unsupervised methods to these data sets, which could result in a different, new classification scheme altogether.

ACKNOWLEDGMENTS

We thank the UKST unit of the Royal Observatory of Edinburgh for the plate material; the APM group at RGO Cambridge for scanning support; M. Irwin, D. Lynden-Bell and S. Maddox for helpful discussions; S. Raychaudhury for giving us the galaxy images; and six experts: R. Buta, H. Corwin, G. de Vaucouleurs, A. Dressler, J. Huchra and S. van den Bergh, for their efforts in classifying the sample. LSJ thanks the financial support provided by FAPESP and CNPq, and the hospitality of the Institute of Astronomy, Cambridge and the Royal Greenwich Observatory. AN acknowledges an Isaac Newton Studentship.

REFERENCES

- Adorf H. M., 1989, in Heck A., Murtagh F., eds, *Knowledge Based Systems in Astronomy*. Springer-Verlag, Heidelberg, p. 215
- Adorf H. M., Meurs E. J. A., 1988, in Seitter W. C., Durebeck H. W., Tacke M., eds, *Large Scale Structure in the Universe: Observational and Analytical Methods*. Springer-Verlag, Heidelberg, p. 315
- Angel J. R. P., Wizinowich P., Lloyd-Hart M., Sandler D., 1990, *Nat*, 348, 221
- de Vaucouleurs G., 1959, in Flüge S., ed., *Handbuch der Physik*, 53, 275, Springer-Verlag, Berlin
- de Vaucouleurs G., 1963, *ApJS*, 8, 31
- de Vaucouleurs G., de Vaucouleurs A., Corwin H. G., Jr., Buta R., Paturel G., Fouqué P., 1991, *Third Reference Catalogue of Bright Galaxies*. Springer-Verlag, New York
- Doi M., Fukugita M., Okamura S., 1993, *MNRAS*, 264, 832
- Dressler A., 1980, *ApJ*, 236, 351
- Hopfield J. J., Tank D. W., 1986, *Sci*, 233, 265
- Hubble E., 1926, *ApJ*, 64, 321
- Hubble E., 1936, *The Realm of Nebulae*. Yale Univ. Press, New Haven
- Lahav O. et al., 1994, *Sci*, 267, 859
- Lauberts A., Valentijn E. A., 1989, *The Surface Photometry Catalogue of the ESO-Upsala Galaxies*. ESO
- Lynden-Bell D., Faber S. M., Burstein D., Davies R. L., Dressler A., Terlevich R. J., Wegner G., 1988, *ApJ*, 326, 19
- McCulloch W. S., Pitts W. H., 1943, *Bull. Math. Biophys.*, 5, 115
- Mackay D., 1992, PhD thesis, California Institute of Technology
- Maddox S. J., Sutherland W. J., Efstathiou G., Loveday J., 1990, *MNRAS*, 243, 692
- Mihalas D., Binney J., 1981, *Galactic Astronomy*, second edition. Freeman, San Francisco
- Morgan M. W., 1958, *PASP*, 70, 364
- Naim A. et al., 1995, *MNRAS*, in press
- Odewahn S. C., Stockwell E. B., Pennington R. L., Humphreys R. M., Zumach W. A., 1991, *AJ*, 103, 318
- Paturel G., Fouqué P., Bottinelli L., Gouguenheim L., 1989, *A&AS*, 80, 299
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1993, *Numerical Recipes*, second edition. Cambridge Univ. Press, Cambridge
- Sandage A. R., 1961, *The Hubble Atlas of Galaxies*. Carnegie Institute of Washington, Washington
- Schweizer F., 1976, *ApJS*, 31, 312
- Simien F., de Vaucouleurs G., 1986, *ApJ*, 302, 564
- Storrie-Lombardi M. C., Lahav O., Sodr  L., Jr., Storrie-Lombardi L. J., 1992, *MNRAS*, 259, 8p
- Tully R. B., Fisher J. R., 1977, *A&A*, 54, 661
- van den Bergh S., 1960a, *ApJ*, 131, 215
- van den Bergh S., 1960b, *ApJ*, 131, 558
- van den Bergh S., 1976, *ApJ*, 206, 883

Galaxy Id	RA (1950)	DEC (1950)	Temp	Sexp	Tann	Sann	Ell	SB	Bulge Slopes	RI	Light	Concentration	Indices	Arms	to Disk	Ratios	#Arms	Leng	Intens
f634g001	9 5 6	-14 16 40.7	0.40	1.1	1.09	0.50	0.53	2.98	5.26	8.91	0.26	0.53	0.84	1.02	0.05	1.58	1.50	17.00	6.70
f634g002	9 4 28.9	-17 27 28.7	0.16	1.1	1.27	0.50	0.57	4.85	1.96	12.90	0.16	0.31	0.46	0.78	0.06	0.29	1.00	10.00	7.71
f634g003	9 4 34.7	-15 26 10.9	5.95	1.67	5.97	0.24	0.51	3.20	1.94	12.70	0.05	0.14	0.60	0.64	0.24	0.34	0.47	17.00	7.57
f634g004	9 4 48.7	-15 6 15.9	5.35	1.47	5.74	0.12	0.54	4.10	0.40	2.38	0.05	0.18	0.32	0.81	0.16	0.69	0.83	11.00	9.00
f634g005	9 4 44.4	-15 13 6.5	9.25	1.22	4.30	0.24	0.66	5.66	1.47	1.60	0.14	0.24	0.39	0.90	0.00	0.00	0.00	0.00	0.00
f634g006	8 48 0.9	-16 13 29.2	4.00	0.82	0.66	0.41	0.30	3.02	3.70	3.66	0.91	0.08	0.92	1.92	0.18	1.17	1.15	7.00	8.86
f634g007	9 17 48.6	-16 18 46.2	4.90	0.69	4.82	0.30	0.45	4.51	2.85	1.83	0.25	0.16	0.29	0.64	0.13	0.56	0.70	23.00	7.43
f634g008	9 13 50.3	-16 6 16.4	3.85	1.46	0.19	0.28	0.73	4.49	7.80	9.37	0.05	0.18	0.37	0.73	0.01	0.14	0.34	2.00	6.00
f634g009	9 29 15.0	-17 20 12.5	3.15	1.07	2.79	0.25	0.44	3.29	4.60	3.53	0.11	0.35	0.96	0.53	0.05	0.36	0.56	1.07	2.00
f634g010	9 29 19.0	-15 49 16.8	4.00	0.00	4.06	0.15	0.47	7.56	2.44	3.17	0.17	0.33	0.50	0.00	0.00	0.00	0.00	0.00	0.00
f634g011	9 28 29.8	-14 30 53.9	6.55	0.80	5.72	0.46	0.87	3.58	3.47	1.09	0.15	0.20	0.32	0.00	0.00	0.00	0.00	0.00	0.00
f634g012	9 26 35.8	-14 35 19.5	4.55	1.17	3.27	0.60	0.67	4.91	0.87	2.13	0.04	0.14	0.26	0.61	0.74	0.86	0.94	3.48	3.48
f634g013	9 24 6.0	-12 29 29.7	5.95	0.99	4.24	0.32	0.50	3.11	7.54	4.70	0.11	0.23	0.27	0.62	0.00	0.00	0.00	0.00	0.00
f634g014	9 13 49.1	-18 10 59.9	1.65	1.26	2.14	0.20	0.76	3.11	0.76	1.68	0.15	0.18	0.32	0.72	0.82	0.91	0.97	3.73	3.73
f634g015	9 18 10.9	-16 15 13.1	0.95	1.07	1.46	0.29	0.57	6.33	6.85	2.01	0.05	0.17	0.30	0.84	0.91	0.97	4.97	4.97	4.97
f634g016	9 32 34.3	-13 26 48.3	3.25	0.94	2.95	0.28	0.31	2.78	4.03	5.00	0.13	0.12	0.24	0.56	0.71	0.84	0.91	12.00	6.58
f634g017	9 40 14.4	-13 28 36.2	3.25	0.76	1.91	0.34	0.54	4.73	6.27	7.88	0.08	0.18	0.35	0.55	0.66	0.75	0.85	0.99	3.27
f634g018	9 40 26.7	-17 21 34.6	5.65	1.46	4.90	0.19	0.60	4.40	2.61	2.17	0.13	0.26	0.73	0.82	0.95	0.99	3.97	4.21	4.21
f634g019	9 48 56.4	-15 48 59.4	4.45	0.75	3.92	0.27	0.34	4.19	0.89	4.87	0.12	0.30	0.44	0.56	0.70	0.81	0.91	9.97	3.55
f634g020	9 43 24.1	-14 8 15.9	1.00	0.00	3.70	0.29	0.47	4.53	9.45	1.16	0.15	0.19	0.38	0.47	0.82	0.94	0.98	4.17	4.17</

APPENDIX A: CLASSIFICATION OF 86 GALAXIES

Table A1 lists 86 galaxies from our sample, giving their identification in the sample, their RA and Dec. coordinates, the classifications they were given, and all the parameters used by the ANNs. The classification fields give the expert type (T_{exp}) and the standard deviation among the expert classifications (S_{exp}) for each galaxy, as well as the ANN mean type (T_{ann}) and its standard deviation over 10 runs

(S_{ann}). The parameters used by the ANNs follow in this order: ellipticity (Ell), surface brightness (SB), bulge size (Bulge), slope ratio (Slopes), Peak rI (rI), the 10 concentration indices (C10, C20, ..., C90, $C_{75/25}$), the six arms-to-disc ratios (AD_{all} , AD_{3-7} , ..., AD_{23-27}), and the arm parameters: number of arms (# Arms), average arm length (Leng.), and average arm intensity (Intens.). All quoted parameters values are as received from the feature extraction software, without normalization.